

# Co-Following on Twitter

Venkata Rama Kiran Garimella  
Qatar Computing Research Institute  
Doha, Qatar  
gvrkirann@gmail.com

Ingmar Weber  
Qatar Computing Research Institute  
Doha, Qatar  
iweber@qf.org.qa

## ABSTRACT

We present an in-depth study of co-following on Twitter based on the observation that two Twitter users whose followers have similar friends are also similar, even though they might not share any direct links or a single mutual follower. We show how this observation contributes to (i) a better understanding of language-agnostic user classification on Twitter, (ii) eliciting opportunities for Computational Social Science, and (iii) improving online marketing by identifying cross-selling opportunities.

We start with a machine learning problem of predicting a user's preference among two alternative choices of Twitter friends. We show that co-following information provides strong signals for diverse classification tasks and that these signals persist even when (i) the most discriminative features are removed and (ii) only relatively "sparse" users with fewer than 152 but more than 43 Twitter friends are considered.

Going beyond mere classification performance optimization, we present applications of our methodology to Computational Social Science. Here we confirm stereotypes such as that the country singer Kenny Chesney (@kennychesney) is more popular among @GOP followers, whereas Lady Gaga (@ladygaga) enjoys more support from @TheDemocrats followers.

In the domain of marketing we give evidence that celebrity endorsement is reflected in co-following and we demonstrate how our methodology can be used to reveal the audience similarities between Apple and Puma and, less obviously, between Nike and Coca-Cola. Concerning a user's popularity we find a statistically significant connection between having a more "average" followership and having more followers than direct rivals. Interestingly, a *larger* audience also seems to be linked to a *less diverse* audience in terms of their co-following.

## Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Sociology; H.3.5 [Online Information Services]: Web-based services; H.2.8 [Database Applications]: Data mining

## General Terms

Experimentation; Human Factors

## Keywords

Twitter; co-following; user similarity

## 1. INTRODUCTION

How much does following a particular set of people reveal about your interests? Does the fact that you follow @Starbucks make it more likely that you follow @TheDemocrats as well? And can Twitter users be grouped in a meaningful way by looking at whether their followers have similar friends<sup>1</sup>?

Such questions are relevant to at least three lines of research. First, there is lots of work on user classification on Twitter, e.g., [8, 9, 3]. Such classifiers often rely on language-specific tools such as stemming or dictionaries with special terms. Our work shows that such information might not be required and a language-agnostic method using a user's friends as features achieves AUC-ROC of .80-.85 for a wide range of binary classification tasks. Second, online social networks are becoming a more and more important data source for Computational Social Science [24, 7]. We contribute to this area by showing how things such as "lifestyle politics" can be studied by using co-following information. Lastly, Twitter with its user base of several hundreds of millions is an important advertising and marketing platform. We show how followership-based similarity methods can be used to identify accounts with a similar audience in terms of interests which could create cross-selling opportunities.

As an example of our approach, consider Figure 1. Here Figure 1a shows the idea that we propose. As an example, consider the twitter accounts of Starbucks (@Starbucks) and the Democrats (@Democrats). Directed edges from a users (user1-user5 in the middle) indicate following behavior of these users. In this example, we see that the two accounts of interest do not share a single common follower. However, many of their followers are "co-following" other, seemingly unrelated accounts. This can be used to deduce the similarity or rather closeness of the two accounts. At first sight, this idea is similar to using common links (co-citations) for clustering web pages [32, 37]. E.g., in Figure 1b, the webpages on the right, CNN and BBC, do not have any links between each other, but could be grouped together using similar links from webpage1 - webpage4. However, typical co-citation or co-linkage approaches would focus on the "1-hop backward" links only and then looking at overlaps.

<sup>1</sup>We use the term "friend" as Twitter terminology referring to another Twitter user that a user follows.

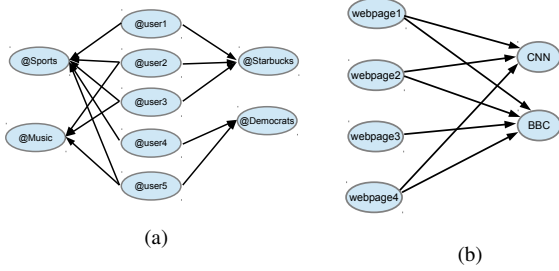


Figure 1: (1a) - Starting from @Starbucks and @Democrats, our co-following analysis uses both the 1st-hop backward edges to the users and the 2nd-hop forward edges to other co-followed accounts. (1b) - Traditional co-linkage clustering of webpages and links would focus on direct 1st-hop backward edges and use shared inlinks to detect similar webpages, in this example CNN and BBC.

In our analysis, we make crucial use of the added “forward” links. In a sense, we are using 2nd order co-citation or co-following rather than ordinary 1st order co-citation.

Most of our analysis is centered around 18 rivalries such as @GOP vs. @TheDemocrats or @McDonalds vs. @BurgerKing. In many cases the two alternatives are arguably interchangeable and one might not expect a big difference between the interests of the followers of, say, @Hertz and @Avis. For each of these seed pairs we obtained up to 2,000 random followers. Their friends are used to construct feature vectors and we perform an in-depth analysis of the co-following behavior. We also construct the same kind of vectors for a set of (i) French and German political parties, and (ii) popular musicians. In all cases the basic hypothesis is that users following similar users are similar and that this propagates to their friends. Concretely, we investigate if users whose followers have similar friends are also similar, even if they might not share any direct followers.

Our findings include the following.

- Using *solely* language-agnostic co-following information provides strong signals concerning a user’s preference even among arguably interchangeable choices such as @Hertz or @Avis.
- Such classification is robust with respect to the removal of the most strongly and often obviously related co-following features.
- Aggregated signals from the general crowd work better for distinguishing binary preferences than relying on the most similar users in a k-NN fashion. For discerning the “center” dual-follower users from the rest the opposite holds.
- Following 50-150 users suffices for a strong signal with little additional gains beyond that.
- A feature analysis confirms stereotypes such that @ladygaga is more popular among @TheDemocrats followers, but also reveals less expected patterns such as that @SnoopDogg followers tend to prefer @Pepsi over @CocaCola.
- There is evidence that celebrity endorsement works as following a celebrity increases the probability of following the related product.

- Users following *both* alternatives of two rivals are sometimes “in the middle” and sometimes a “class by themselves”, thus making it harder to detect such users than either of the two poles.
- Groups of related Twitter accounts, such as musicians or political parties, can be mapped in a simple manner by looking at their followers’ friends. Such a mapping reveals interesting facts, for example, that Apple and Puma target a similar, metropolitan audience.
- There is statistically significant evidence that Twitter users whose followership is closer to a global “average” have more followers than their rivals.
- There is a positive effect on the total follower count by having more intrinsically similar and coherent followers.

To the best of our knowledge, this is the first such study of co-following on Twitter. We hope that both our analysis and our tools will be of interest to researchers working on user classification, Computational Social Science, or on Social Media marketing.

The rest of this paper is organized as follows. First we review related work from domains such as Twitter user classification or follower recommendation. Our data set and pre-processing steps are described in the following section, Section 3. In Section 4 we investigate the potential of co-following information for discerning preferences among similar but rivaling accounts. In Section 4.3 we show how co-following information can be used to group accounts with similar followership in a meaningful manner. In Section 5 we look at how co-following based similarity can be used to visually discriminate between various accounts from domains as diverse as politics and music. Finally, we look at whether it pays off to have mainstream followers in Section 6 before concluding in Section 7.

## 2. RELATED WORK

A key assumption to group users based on the similarity of their followers’ friends is that following is an expression of topical interests or demographic similarity, rather than personal contacts. Kwak et al. address the question whether Twitter is more of a social network or a news media [22]. Looking at trending topics they find “that the majority (over 85%) of topics are headline news or persistent news in nature”. They also observe a low degree of reciprocity among links, which is atypical for normal social networks. Note that for our applications it is *good* that links on Twitter are not dominated by real social ties as this indicates that they are more related to interests. An early study of why people use Twitter is presented in [20]. The authors find “that people use microblogging to talk about their daily activities and to seek or share information”. Given the findings of Kwak et al., the information sharing fraction has likely increased since which, again, benefits our approach. A similar conclusion “that Twitter users have a very small number of friends compared to the number of followers and followees they declare” is reached in [18], further indicating that the majority of follow links is likely related to topical interests, not real-life personal connections. This observation is also confirmed in [38] where Wei et al. find that mutual following is a signal for real life friendship.

Related to our approach of using co-following to measure user similarity is the work in [2] where An et al. use the common audience of two Twitter accounts as a measure of closeness. This differs fundamentally from our methodology as we use *second order*

co-following. Concretely, two accounts that do not share a single follower are considered similar by us if their followers share many friends. We believe that such an approach is preferable to break out of the “homophily ghetto” where users follow what their friends follow. It also allows for a more far-reaching notion of similarity that could be used to, say, align political parties in different countries on a common spectrum, even if no user follows parties in two distinct countries. The benefits of using higher-order co-occurrences rather than only direct co-occurrences have been shown for document retrieval via dimensionality reduction [4]. Outside the scope of Twitter, authors have also studied first order co-citation and co-linking patterns. Small introduces the idea of looking at co-cited academic papers in [33] with the aim of “provid[ing] a new way to study the specialty structure of science”. Culnan looks at author similarity in the realm of Management Information Systems by looking at which authors get co-cited [10]. Common, direct inlinks are used as one of the signals in [35] to form the basis of webpage clustering.

Several papers have looked at various Twitter user classification tasks, typically for (i) political orientation in the US, (ii) gender, and (iii) age [27, 30, 41, 8, 9, 3]. This line of work usually involves a broad set of features, including textual content, network and activity based features, as well as a variety of classification approaches that make use of label-propagation across social links. Our approach in Section 4 differs in a number of ways. First, the binary classification task of “does the user follow A or B” is different. Second, we do not use any content-based features. Third, we do not use any retweet signals as we are not interested in the sparse network of interests that users strongly engage with, but rather the larger network of “weak interests”. Fourth, we do not make use of label-propagation across social links as we are not interested in methods that work for (and re-inforce) “information bubbles” but we are looking for approaches that can be transferred to completely new domains where users do not yet have any direct social ties. Finally, the actual classification performance is of less interest to us than an understanding of how much information is contained in co-following and how this could be used for different applications.

Given its global popularity and the relative ease of data acquisition through open APIs<sup>2</sup>, Twitter has been used before for studies in Computational Social Science and some examples are discussed in the following. Quercia et al. use a crawl of 228K Twitter profiles to test “whether established sociological theories of real-life networks hold in Twitter” [29]. They find that, “much like individuals in real-life communities, social brokers [...] are opinion leaders who tweet about diverse topics, have geographically wide networks, and express not only positive but also negative emotions”. Garcia-Gavilanes et al. link cultural variables such as Pace of Life, Individualism and power to tweeting behavior from several countries [15]. They find strong correlations between these variables and online behavior. Though our current paper only partly falls under the umbrella of Computational Social Science, we hope that this work still illustrates the potential of using co-following information for such studies.

Our mapping and visualization of similar accounts in Section 5 is conceptually similar to community detection which also identifies groups of related accounts. Classical community detection looks at the global network and tries to find areas with unusually high triadic closure [25] or otherwise unusual linkage patterns. A sur-

vey of the area and experimental comparisons can be found in [13] and [23] respectively. Related are notions of graph centrality that also require the whole graph [14]. The approach in [26] is arguably the most similar to ours as it combines local structural information with node similarity in an iterative manner. The similarity used, however, requires direct co-links and even direct links. Our approach differs in a number of ways. First, we are not interested in clustering/grouping *all* users but only understand the relative positions of main ones. Second, we do not require a global view of the network but, in return, we cannot argue about things such as node centrality or PageRank. Third, we do not want to find communities induced by friends-of-friends type links. Rather we strive for a similarity-only based approach that can easily be transferred to domains without any friends-of-friend links.

Determining which of two alternatives a Twitter user is more likely to follow is related to friend recommendation or link prediction as, in a sense, we are suggesting which of the two links should be formed. Intuitively, transitivity and mutuality of links are important signals for link prediction [17] but, as discussed previously, we do not want to use such “three people you follow also follow X” information as it leads to a different type of application, closer to community detection. User similarity based on user attributes has also been used as a feature for link prediction [40, 39, 19]. But this work still partly relies on mutuality and transitivity, which is equivalent to re-enforcing partisan camps without noting any existing similarities in terms of shared interests. For example, such approaches would most likely fail to pick up the similarity between @Puma and @TheAppleInc that we observe in Section 5.3.

Related to our analysis in Section 5 on mapping and visualization of similar users is previous work on segmenting online user populations as our methods can also be applied for audience analysis on Twitter. For web search a clustering based market segmentation is presented in [36]. The authors include demographic variables in their analysis which could also be obtained for Twitter through the help of machine learned classifiers. For web browsing a similar study, also including demographic variables, is presented in [16]. In our work we do not explicitly segment millions of Twitter users but, rather, use regular Twitter users to segment a small number of important accounts.

As mentioned in Section 1, co-citations and common links have been used as a measure of similarity of scientific articles and webpages for a long time [32, 37]. The idea there is that if two webpages have similar inlinks (think “endorsing followers”), they are similar. Our idea is similar to these in the sense that we are using the notion of “backward” links (pages linking to = authors citing = users following) to look for similarity, but it differs as (i) we not only make use of the 1st-hop backward links, but also use 2nd-hop forward links (other accounts followed by a user) to measure similarity. Finally, our application scenario is different as we are not interested in finding, say, topically similar web pages but, rather, trying to find hidden closeness between a priori unrelated accounts on Twitter.

### 3. DATA

Our data set is constructed around a set of *Twitter seed accounts*. These accounts correspond to “rivalries” between two entities such as @CocaCola vs. @Pepsi or @Samsung vs. @TheAppleInc. The full list of 18 account pairs can be found in Table 1. The list of these

<sup>2</sup><https://dev.twitter.com/>

rivalries was obtained from<sup>3</sup>. We only considered those rivalries that are relevant and had a twitter account. To study the “middle” users in Section 4.3, we randomly sampled a set of 2,000 users who follow *both* rivaling accounts. Later, we also look at *groups* of seed accounts, namely, Twitter accounts for (i) German and French political parties, (Section 5.1), (ii) popular musicians, (Section 5.2), and (iii) all the 18 rivalries combined, (Section 5.3). In all cases, we first obtained a list of all the accounts’ followers. From this list we then sampled uniformly at random a set of 2,000 followers. For the cases of rivalries, we also imposed the constraints that the followers were located in the United States. This was done to avoid picking up differences in international market penetrations, rather than within-US cultural differences. The followers for the political parties from Germany and France were also limited to Germany and France respectively. To implement the within-country restriction, we first sampled a set of 10,000 followers for each seed user uniformly at random (or using all when the account had fewer followers), then obtained all of these users’ Twitter bios and self-declared locations. For users with a non-empty profile location string, we ran them through Yahoo! Placemaker<sup>4</sup> to map their locations to countries. We then kept 2,000 users with the appropriate country.

For each of the sampled followers we obtained the full list of their Twitter friends, i.e., users that they follow. The seed accounts pertaining to the corresponding rivalry/group were removed from these friends lists and the remaining ones were treated as a feature vector with each dimension corresponding to a Twitter account being followed. Users who followed *only* seed accounts were dropped.

We used this data to construct a binary classifier for which we created train and test splits, each consisting of  $\sim 1,000$  users. Note that even though 2,000 users were sampled, due to limitations in the Twitter API, the actual number of users for which we could get the friends varies between 1,800-2,000. This could be due to changes in users’ privacy settings, accounts getting blocked and so on. For constructing the training vectors, we only considered users who were followed by at least two users in our training set. That is, if only one of the thousands of users in the training set followed @phdcomics, then following @phdcomics would not be used as a feature. This serves as a simple method for reducing the dimensionality as well as removing unimportant dimensions. This is analogous to the text mining scenario of removing rare tokens with a frequency 1.

## 4. CO-FOLLOWING AND BINARY PREFERENCES

In this section we look at how much a user’s choice of Twitter friends reveals about their preference among two alternatives such as @CocaCola vs. @Pepsi. We do this with different research questions in mind. First, we approach things from a machine learning perspective with an evaluation of the corresponding binary classification task. Next, we do a feature analysis to see which arguably irrelevant features, such as musicians followed, provide information about a user’s soft drink or political preferences. Finally, in Section 4.3, we look at whether we can identify “the middle” or “undecided users” who follow *both* alternatives. Identifying this users segment is of particular relevance for political advertising.

<sup>3</sup><http://money.cnn.com/gallery/news/companies/2013/03/21/greatest-business-rivalries.fortune/index.html>

<sup>4</sup><http://developer.yahoo.com/boss/geo/>

*Feature Vectors Using IDF.* As a preprocessing step, we transformed our binary X-follows-Y vectors to an IDF-weighted alternative. To illustrate why, imagine that almost everybody follows @SuperCelebrity. Then following @SuperCelebrity is not very informative or discriminative and is given a very low IDF weight. For each of the 18 rivalries, we compute the IDF scores of the friends of the followers of the 36 seed rivals. In total, 63,853 (N) followers of the seed accounts were used ( $= 36 \times 2,000$ , minus cases with fewer than 2,000 followers and blocked/deleted/private accounts). We computed the IDF of each of their friends, obtaining one IDF-weighted vector for each of the 63,853 followers.

$$IDF(user_i) = \log \left( \frac{N}{|followers(user_i)|} \right), \quad (1)$$

where  $followers(user_i)$  indicates the followers of a particular user, from the set of followers sampled for the seed rival accounts. Each of these IDF-weighted vectors was then normalized in 2-norm and, for a given seed account, all of its followers normalized vectors are summed up. This summed vector is then re-normalized in 2-norm to give the final “global” summary vector for the seed account.

### 4.1 Machine Learning Performance

In this section we evaluate how much information co-following provides for the task of classifying users according to their binary preference. The ground truth is the single<sup>5</sup> account that the test user actually followed. The feature dimensions corresponding to the following of the seed accounts are always removed and, later, we also remove strongly correlated features such as following @BarackObama for the @GOP vs. @TheDemocrats task. Empty vectors, after removing the seed accounts (for users following only the seed accounts) are ignored. Our main performance measure is the area-under-curve (AUC) for the Response Operating Characteristic Curve (ROC) as computed by<sup>6</sup> [11]. We also report AUC for the Precision-Recall Curve (PR) though AUC-ROC will be the default. A value of 0.5 indicates a random, unskilled prediction model.

Note that we are more interested in understanding the relative performance when, say, the most discriminative features are removed than we are in achieving the highest possible classification accuracy. The accuracy could always be improved further by using other algorithms (SVM, Maximum Entropy, etc.), other feature sets (textual data, interaction features, network features, etc.), or incorporating other techniques (label propagation, community detection, etc.). Our focus is more on understanding issues such as robustness under feature removal, relative performance on sparse test vectors or opportunities for Computational Social Science arising from feature analysis.

*Global vs. Local Approach.* To determine whether you fall into group A or B, is it more useful to know (i) what the general, average members of A and B are like, or (ii) which of any of the two contains a small number of members just like you? The answer to this question has applications both for the design of classification algorithms and for understanding the structure of groups of followers. We try to answer this question by comparing two

<sup>5</sup>We will later in Section 4.3 look at those users individually that follow *both* alternatives.

<sup>6</sup><http://mark.goadrich.com/programs/AUC/>

different classification strategies. First, a “global” method using the single IDF-weighted summary vector described above. This method also includes information about fairly rare friends as it aggregates information from about 2,000 followers. Second, a “local” approach, that uses a k-nearest neighbor classifier. It then assigns each test vector to the class with the largest number of close neighbors among the top k. For k-NN we experimented with a range of values for k from 1 to 9 in increments of 2. There was a clear tendency for higher values of k to perform better and so we stuck to a choice of k=9. We did not experiment with larger values as the general trend of a more and more global approach performing better was our main objective, rather than identifying an optimal value of, say, k=135.

The performance of the binary classification is shown in Table 1. The global approach always performs better than the local approach, showing that it is worth aggregating the long tail of rare co-follower relations. The AUC-ROC averaged across the 18 tasks is 0.81.

Rivalry	Global	Local
@Budweiser vs. @MillerCoors	0.86 (0.91)	0.80 (0.85)
@FedEx vs. @UPS	0.73 (0.73)	0.69 (0.72)
@GM vs. @Ford	0.75 (0.86)	0.69 (0.76)
@GOP vs. @TheDemocrats	0.91 (0.95)	0.86 (0.93)
@Hertz vs. @Avis	0.92 (0.93)	0.91 (0.92)
@InsideFerrari vs. @lamborghini	0.92 (0.95)	0.87 (0.93)
@jcpenny vs. @Sears	0.75 (0.82)	0.67 (0.72)
@McDonalds vs. @BurgerKing	0.78 (0.79)	0.68 (0.70)
@MercedesBenz vs. @bmw	0.89 (0.93)	0.86 (0.91)
@Nike vs. @Reebok	0.78 (0.74)	0.73 (0.68)
@NikonUSA vs. @CanonUSAimaging	0.83 (0.85)	0.78 (0.83)
@pepsi vs. @CocaCola	0.69 (0.76)	0.65 (0.73)
@PUMA vs. @adidas	0.77 (0.84)	0.69 (0.73)
@SamsungMobile vs. @TheAppleInc	0.95 (0.96)	0.92 (0.94)
@Starbucks vs. @DunkinDonuts	0.80 (0.87)	0.72 (0.82)
@Target vs. @Walmart	0.78 (0.86)	0.69 (0.79)
@thewanted vs. @onedirection	0.79 (0.88)	0.76 (0.84)
@Visa vs. @MasterCard	0.71 (0.72)	0.62 (0.59)

Table 1: Performance comparison for the 18 binary classification tasks (detecting preference among rivaling alternatives) in terms of AUC-ROC (AUC-PR) for both the global and local similarity-based approaches.

Note that the performance is arguably bounded by the fraction of fake accounts as such accounts might not obey any meaningful co-following behavior. According to <sup>7</sup>, the typical fraction of fake followers for our 36 accounts is 20-30% and it goes up to 50%/39% for @CocaCola/@pepsi, the worst-performing pair in our set, whereas it is only 5%/12% for @Hertz/@Avis with one of the best performances. However, for @FedEx/@UPS the classification performance is also comparatively poor, despite the low fake follower fractions of 12%/13%. This indicates that fake followership is not the only element influencing classification performance and that certain tasks are intrinsically harder.

*Removing Obvious Co-following Signals.* Discovering that following @BarackObama on Twitter is an indication for following @TheDemocrats rather than @GOP is obvious. Similarly, following @CokeZero correlates positively with following @CocaCola.

<sup>7</sup><http://fakers.statuspeople.com/Fakers/Scores>

As we were more interested in studying the *non-obvious* dependencies we investigated the classification performance when the most predictive features are removed. Note that this is the *opposite* of what normal feature selection does.

Concretely, for each binary setting we rank features as follows. For each rivalry pair A, B, we look at the absolute differences in the feature values listed in A’s and B’s summary vectors. These absolute differences are then sorted in descending order. Features with a large difference correspond to accounts that are typically more followed by the followers of one seed account, but not the other one. Pseudo-code to compute the difference is shown in Algorithm 1. We then rank these features by the absolute difference.

**Algorithm 1** Pseudocode for computing difference between features

---

```

for feature  $\in$  train_A.features do
  if feature  $\in$  train_B.features then
    diff  $\leftarrow$  train_A[feature] – train_B[feature]
  else
    diff  $\leftarrow$  train_A[feature]
  end if
end for
for feature  $\in$  train_B.features do
  if feature  $\notin$  train_A.features then
    diff  $\leftarrow$  –train_B[feature]
  end if
end for
return diff

```

---

In order to check the influence of the top features, we removed the top 10, 20, 50, 100 and 200 most obvious features and compared the AUC in each case. Note that in this setting, since we remove the most influential features, the size of the test set might change (because some users might only follow these influential users). In order to compensate for this, we tried two variants, one considering only users who have more than 201 followers, so that the size of the test set is fixed and the other with a varying test set size. The results of the former case are presented in Figure 2, though results in both cases are comparable. The y-axis indicates AUC averaged across all the rival groups. We note a gradual decrease in the mean AUC as we remove more features, which is in line with what is expected.

For our later “mapping” analysis (Section 5), we take a similar approach to remove the top 20 features for each of the A vs. non-A classification problems, where A iterates over all the seed accounts and the non-A group pools all non-A seeds.

*Influence of Number of People Followed.* Different users reveal different amounts of their preferences on Twitter. Some choose to follow hundreds of users and some follow none. How much following information is needed to gauge a user’s preference? Does the number of users you follow affect the performance of the classification? We address these questions here by bucketing the users in the test set according to the number of their Twitter friends. Apart from the seed users themselves, the 20 most discriminative (and obvious) features were ignored (see previous discussion) and only users following at least one other user were considered. We bucketed the test users according to the number of friends they have. In order to make qualitative conclusions over the general population, we computed these buckets by pooling the followers of all seed users. We then found the 25<sup>th</sup>, 50<sup>th</sup>, 75<sup>th</sup> and 100<sup>th</sup> percentiles for the number of friends these users have and bucketed the users




Figure 2: Average AUC-ROC across the 18 binary classification tasks (detecting preference among rivaling alternatives) as more and more features are removed, but the test set size is fixed. Only users with more than 201 followers are considered. Error bars indicate the standard error across the tasks.

for each group accordingly. The buckets were  $bucket1 \in [1, 43)$ ,  $bucket2 \in [43, 152)$ ,  $bucket3 \in [152, 446)$ , and  $bucket4 \in [446, 1, 363, 508]$ . Figure 3 plots the average AUC across the four user buckets. As expected, the more followers you have, the easier it is to know your relative preference. An interesting observation is that for the “global” approach, the curve almost subsides after the first bucket, indicating that the value added after 43 followers is not noticeable. It also indicates that  $\sim 100$  users is sufficient to predict the class of a user with  $> 0.8$  AUC-ROC.

## 4.2 Feature Analysis

*Correlation with Lifestyle.* Apart from using it for user classification and target advertising, co-following patterns are also of interest in their own right and can serve to answer questions in Computational Social Science. For instance, there is academic work that looks at “lifestyle politics” such as the correlation between political leaning and television preferences or the stereotype that liberals like lattes [6, 5, 28]. Our approach contributes to this by offering a language-agnostic tool to use online data to quantify such effects at scale. In this section, we use the feature ranking described previously to generate the top, discriminative features (such as following @BarackObama to predict a @GOP or @TheDemocrats preference).

As an example, we look at the rivalries between @GOP vs. @TheDemocrats and @Pepsi vs. @Cocacola. For both cases we look at the top discriminative co-following features from the <http://WeFollow.com> categories Music, Sports and News<sup>8</sup>. Table 2 shows some examples of the insights we can get from co-following patterns. The lifestyle correlations for the political rivalry @GOP vs. @TheDemocrats can be inspected to make intuitive sense with, e.g., @nytimes being more popular among @TheDemocrats followers.<sup>9</sup> For @Pepsi vs.

<sup>8</sup>A small number of mislabeled entries were removed. For example, @AgainAmerica was incorrectly listed in the Music category.

<sup>9</sup>The New York Times is generally perceived to have a liberal bias, see [http://en.wikipedia.org/wiki/The\\_New\\_York\\_Times#Political\\_persuasion\\_overall](http://en.wikipedia.org/wiki/The_New_York_Times#Political_persuasion_overall).




Figure 3: Average AUC-ROC across the 18 binary classification tasks (detecting preference among rivaling alternatives) for different sets of users. Though the AUC increases for users with more Twitter friends, the growth slows down already when the number of users crosses 43, when the global similarity based approach is used.

@CocaCola many observations can be explained by the fact that @Pepsi targets the younger “New Generation”. In passing we also mention that following @Starbucks is indeed an indicator for co-following @TheDemocrats rather than @GOP.

*Celebrity Endorsements.* An interesting side note of examining the top features is the detection of celebrity endorsements. By observing these features, we found out that celebrity endorsements go hand in hand with who people follow. Some examples include Derrick Rose (@drose) for Adidas; Kevin Durant (@KDTrey5) for Nike; SnoopDogg (@SnoopDogg), Nicki Minaj (@NickiMinaj) and Drake (@drake) for Pepsi, and Maroon 5 (@maroon5) and David Guetta (@davidguetta) for Cocacola, and not so famous relations like Dan Bailey (@Dan\_Bailey9) and Rich Froning (@richfroning) to Reebok. (They are both stars of the show CrossFit Games, hosted by Reebok.) This observation has interesting applications in marketing campaigns and recommendation systems and deserves more analysis in the future.

## 4.3 Analyzing the “Center”

For certain marketing campaigns, in particular in the political domain, identifying strong supporters or sworn enemies is less useful than identifying users in the likely *center*. In this study, we define the center as users that follow *both* of two alternative and rivaling accounts. We investigate whether this center is then indeed “in the middle” and whether such users can be told apart from the more polarized users.

*Is the Center in the Middle?* One definition to lie “in the middle” is that the distance to either end of the spectrum should be smaller than the distance from one end of the spectrum to the other. If we think of the triangle formed by the two alternatives  $A$  and  $B$  and the supposed center  $C$  then we require that both of the lengths  $|AC|$  and  $|CB|$  are smaller than  $|AB|$ . In our case, we are looking at 2-normalized vectors lying on the unit sphere and so we require



	@GOP	@TheDemocrats
Music 1	@kennychesney (64)	@ladygaga (24)
Music 2	@jakeowen (122)	@aliciakeys (57)
Music 3	@taylorswift13 (139)	@SnoopDogg (62)
Sports 1	@espn (125)	@rolandsmartin (66)
Sports 2	@runnersworld (143)	@bubbawatson (79)
Sports 3	@AdamSchefter (178)	@NBA (188)
News 1	@WSJ (12)	@nytimes (11)
News 2	@HumanEvents (77)	@cnnbrk (21)
News 3	@toddstarnes (107)	@NYTimeskrugman (23)

---

	@Pepsi	@Cocacola
Music 1	@SnoopDogg (1)	@maroon5 (8)
Music 2	@Nickiminaj (2)	@davidguetta (19)
Music 3	@Drake (5)	@Pitbull (28)
Sports 1	@shaq (3)	@SInow (99)
Sports 2	@ochocinco (20)	@kaka (103)
Sports 3	@DwightHoward (42)	@chicagobulls (206)
News 1	@Rapup (120)	@cnnbrk (55)
News 2	@Life (133)	@WSJ (81)
News 3	@MTVnews (339)	@TheOnion (183)

Table 2: List of differentiating co-following features from different WeFollow classes, for @GOP vs. @TheDemocrats, and @Pepsi vs. @Cocacola. The numbers in parentheses indicate the absolute position of this feature in our ranking irrespective of the topic (Music, Sports or News).

that both  $b = \angle(A, C) < c = \angle(A, B)$  and  $a = \angle(C, B) < c = \angle(A, B)$ . Figure 4 shows an illustration where  $A$  and  $B$  refer to the unit vectors of the centroids of the two corresponding rivals, with  $C$  referring to the centroid of the dual-followership. The results of this test of the angles which we refer to as  $\Delta$  is presented in Table 3. For 12 of the 18 cases this condition is satisfied, but for six cases it is violated. This gives evidence that in general users following two similar but rivaling alternatives are indeed *between* the two ends, but due to the small number of rivalries considered the result is not statistically significant. In future work we plan to look more at when this fails to hold and when these users are not between but rather *different*. To quantify how much in the middle




Figure 4: Illustration of basic spherical trigonometry that is used to quantify the degree of centered-ness of the dual-followership crowd (centroid  $C$ ) with respect to the two alternatives' centroids ( $A$  and  $B$ ).

the dual-followership is, we also look at the angle of the spherical triangle formed at  $C$ . See Figure 4 for an illustration. The larger the angle at  $C$  the closer to the middle  $C$  is. This angle is listed in Table 3 and it indicates that for @pepsi vs. @CocaCola  $C$  is furthest away from the middle whereas for @Hertz vs. @Avis it is closest. Note that in no case the angle formed at  $C$  exceeds  $90^\circ$  indicating that  $C$  is both partly a “bridge” in the middle but still largely a class by itself. We will explore this in detail in the next paragraph.

Rivalry	$\Delta$	angle at $C^\circ$
@Budweiser vs. @MillerCoors	Yes	72
@FedEx vs. @UPS	Yes	65
@GM vs. @Ford	Yes	66
@GOP vs. @TheDemocrats	Yes	64
@Hertz vs. @Avis	Yes	76
@InsideFerrari vs. @lamborghini	Yes	65
@jcpenney vs. @Sears	Yes	68
@McDonalds vs. @BurgerKing	No	64
@MercedesBenz vs. @bmw	Yes	68
@Nike vs. @Reebok	No	64
@NikonUSA vs. @CanonUSAimaging	No	68
@pepsi vs. @CocaCola	No	58
@PUMA vs. @adidas	Yes	64
@SamsungMobile vs. @TheAppleInc	Yes	66
@Starbucks vs. @DunkinDonuts	Yes	66
@Target vs. @Walmart	Yes	66
@thewanted vs. @onedirection	No	66
@Visa vs. @MasterCard	No	62

Table 3: Results of whether users following *both* alternatives can be considered to lie “in the middle”.  $\Delta$  refers to the necessary condition of the angles, see text.

*Can We Tell the Middle From the Poles?* In the previous paragraph we saw both evidence for the dual-followership crowd *lying between* the two ends (as  $\Delta$  was satisfied in 12 out of 18 cases) but also for being a *qualitatively different* class of their own (as the angle at  $C$  never approached  $90^\circ$ ). In this paragraph we look at whether the dual-followership can be told apart from the two poles and which features are of relevance in this process. For the classification, we compute the three centroids for the two rivals and dual-followership following the same procedure as before. Then for the  $\sim 2,000$  test instances, roughly 500 from each pole and 1,000 from the dual group, we looked at with which centroid they shared the highest similarity. For Table 4 we look only at the Both-or-Not performance where the two rivals would be treated as one class.

Looking more closely at the relative closeness of the three classes, we also looked at the three-class classification problem and the corresponding confusion matrices. Concretely, we looked at the number of rival pairs A-B for which both  $\#(\text{true A, closest to C}) > \#(\text{true A, closest to B})$  and  $\#(\text{true B, closest to C}) > \#(\text{true B, closest to A})$ , where C corresponds to the dual-followership. Intuitively, if both of these conditions are satisfied for a pair then C is the most confusing class, indicating it is between the others. Surprisingly, for only three out of the 18 cases this was the case. The story of when dual-followership indicates a “middle ground” therefore seems to be complex.

## 5. MAPPING THE TWITTERSPHERE VIA CO-FOLLOWING

Rivalry	Global	Local
@Budweiser vs. @MillerCoors	0.88 (0.91)	0.87 (0.85)
@FedEx vs. @UPS	0.79 (0.74)	0.81 (0.81)
@GM vs. @Ford	0.73 (0.79)	0.77 (0.79)
@GOP vs. @TheDemocrats	0.78 (0.72)	0.89 (0.86)
@Hertz vs. @Avis	0.83 (0.91)	0.77 (0.84)
@InsideFerrari vs. @lamborghini	0.79 (0.78)	0.84 (0.82)
@jcpenny vs. @Sears	0.78 (0.76)	0.81 (0.81)
@McDonalds vs. @BurgerKing	0.77 (0.74)	0.85 (0.85)
@MercedesBenz vs. @bmw	0.73 (0.73)	0.80 (0.78)
@Nike vs. @Reebok	0.77 (0.69)	0.85 (0.82)
@NikonUSA vs. @CanonUSAimaging	0.79 (0.75)	0.84 (0.83)
@pepsi vs. @CocaCola	0.84 (0.79)	0.90 (0.87)
@PUMA vs. @adidas	0.91 (0.91)	0.91 (0.88)
@SamsungMobile vs. @TheAppleInc	0.78 (0.74)	0.86 (0.84)
@Starbucks vs. @DunkinDonuts	0.81 (0.84)	0.82 (0.80)
@Target vs. @Walmart	0.76 (0.78)	0.78 (0.77)
@thewanted vs. @onedirection	0.82 (0.85)	0.76 (0.73)
@Visa vs. @MasterCard	0.77 (0.81)	0.77 (0.81)

Table 4: Performance comparison for the 18 binary classification tasks (detecting the “center” from the two poles) in terms of AUC-ROC (AUC-PR) for both the global and local similarity-based approaches.

In this section we look at whether a co-following based similarity can be used to map the relative positions of players from domains as diverse as politics and music. Though our maps can be seen as “community detection”, the approach and interpretation is very different. Whereas traditional community detection algorithms use direct social links and, e.g., would try to find clusters with unusually high triadic closure [25], our approach relies on more indirect and high-order links. As an example, imagine two football clubs that are fierce rivals and who would definitely not follow each other. Fans and Twitter followers of either club might also not follow the other one. However, their fans might jointly follow many other accounts related to sports news. Due to this co-following of the clubs’ followers we would consider the clubs as similar in terms of their audiences’ interests. This approach also opens up opportunities for cross-marketing and cross-selling: if two Twitter accounts from different domains share a similar followership then they might consider cross-posting or otherwise combining their forces. Note again that they do not have to share even a single follower to be considered similar as we look at second-order following relations, namely, the friends of their followers.

Technically, we did the following: For each of the  $\sim 2,000$  followers of an account, we constructed the IDF vector for the users they follow.<sup>10</sup> We then computed the pair-wise cosine similarity between these feature vectors. Since we need distances, we used (1 - cosine similarity) as the measure of distance. We then used the classical, Metric Multi-Dimensional Scaling (MDS) [34] on this data with the `cmdscale()` function in *Matlab*.

Note that MDS is a *lossy* embedding and that even though two points appear close in the 2-dimensional plane, they might be far apart in the original high dimensional space. Therefore, all conclusions and observations we derived from such mappings in the following have also been validated using the high dimensional similarity information.

<sup>10</sup>The total number of followers (N) was 38,358 (Musicians), and 7,783 (Germany), and 9,966 (France).




Figure 5: A 2D MDS similarity map of French and German political parties. Similarity measures are derived from their followers’ aggregated friends.

## 5.1 Political Parties in Germany and France

Both for France and for Germany we obtained a list of parties/alliances that were present in the corresponding national parliaments with a non-trivial number of seats. We were interested to see if the grassroots-derived party similarities matched the ideological realities. For France our list included: Europe Écologie Les Verts/Europe Ecology The Greens (@EELV, center-left, 104,997 followers), Front de Gauche/Left Front (@FDG, left, 17,481 followers), Union des Démocrates et Indépendants/Union of Democrats and Independents (@UDI\_off, center, 7,825 followers), Union pour un Mouvement Populaire/Union for a Popular Movement (@ump, center-right, 75,292 followers), Parti socialiste/Socialist Party (@partisocialiste, center-left, 67,693 followers), and Front national/National Front (@FN\_officiel, far right, 25,429 followers).

For Germany we used: Die Linke/The Left (@dieLinke, left, 22,746 followers), Freie Demokratische Partei/Free Democratic Party (@fdp\_de, center-right, 17,006 followers), Die Grünen/The Greens (@Die\_Gruenen, center-left, 72,920 followers), Christlich Demokratische Union & Christlich Soziale Union/Christian Democratic Union & Christian Social Union (@cdusubt, center-right, 19,998 followers), Sozialdemokratische Partei Deutschlands/Social Democratic Party of Germany (@spdde, center-left, 45,407 followers).

For both countries the results make intuitive sense. For France (Figure 5, left) the Front de Gauche (@fdg) and the Front national (@fn\_officiel) are the two parties furthest away from the center in terms of political spectrum, whereas the other center-left and center-right parties are close to each other. Similarly, for Germany (Figure 5, right), the isolated outlier is Die Linke (@dielinke), a party which is frequently boycotted/ignored by other parties in terms of coalition considerations, and the remaining parties are close to their respective frequent coalition partners, namely, Die Grünen (@die\_gruenen) & Sozialdemokratische Partei Deutschland (@spdde); Freie Demokratische Partei (@fdp\_de) & Christlich Demokratische/Soziale Union (@cdusubt). Note that it is not obvious that the similarity of parties’ followers should agree with the parties’ orientation as the majority of a follower’s Twitter friends will be apolitical. Also see Section 4.2 for examples of “lifestyle politics” in the case of the US.

## 5.2 Popular Musicians

To see if our approach generalizes to other domains, we decided to map popular musicians on Twitter. To this end we obtained a list of the top 22 musicians (in terms of prominence score) from <http://wefollow.com/interest/music>.<sup>11</sup>

Figure 6 shows the map that was obtained in the same way as as the previous case. Most of the observed structure corresponds

<sup>11</sup>We removed 4 accounts from the initial list that corresponded to media/producers rather than musicians or bands.






Figure 6: A 2D MDS similarity map of popular musicians. Similarity measures are derived from their followers’ aggregated friends.

to musical genres. For example, Lil Wayne (@liltunechi), Chris Brown (@chrisbrown) and Drake (@drake) are rappers and are co-mapped together in the map, marked in red. Similar is the case of Snoop Dogg (@snoopdogg) and Kanye West (@kanyewest), marked in green, both of which are hip hop artists. However, there are also surprising things that emerge such as the relative closeness of “Weird Al” Yankovic (@alyankovic), famous for musical parody, and Yoko Ono (@yokoono), both marked in orange. Though very different musical genres, both arguably appeal to an older, more educated audience. This already hints at applications of such analysis for the identification of cross-selling opportunities.

### 5.3 Combination of All Rivals

To show the full generalizability of this mapping approach also *across* domains we combined all the 36 Twitter accounts from the 18 rivalries and mapped them in a common space in Figure 7. As one would expect, many rivals such as @Target vs. @jcpenney and @thewanted vs. @onedirection are comparatively close as their followers share similar interests. However, the relative distances across rivalries also makes sense. For example, the beer brand @MillerCoors is closer to @GOP than to @TheDemocrats and the opposite holds for @Budweiser. This makes sense as it has been observed before that “Republicans are also big fans of Miller Lite and Coors Light, but Democrats drink more Budweiser” [12, 31], though, some studies show that the opposite is true [1]. Sometimes, studies such as this one are inconclusive and show conflicting results based on the demographics studied (such as voters vs. just politically leaning but not necessarily voting), sampling methods used, etc. Similarly, @GM and @Avis are very close in the low-dimensional embedding. Again, this makes sense as “[s]ince the late 1970s, Avis has featured mainly General Motors (GM) vehicles”<sup>12</sup>. Also noteworthy is the closeness between @PUMA and @TheAppleInc. Though there does not appear to be any formal alliance, both brands try to create a similar image of themselves. Puma targets the “sports lifestyle” trend with persona attributes such as metropolitan and international [21] which, arguable also apply to Apple.

We believe that such a mapping is useful to quickly generate hypotheses for lifestyle politics and similar research areas that can then be investigated in depth. It is important to note that even if some of these findings were not to hold “offline” in all cases, these Twitter-only findings are still useful for online advertising as they definitely provide a signal.

<sup>12</sup>[http://en.wikipedia.org/wiki/Avis\\_Rent\\_a\\_Car\\_System](http://en.wikipedia.org/wiki/Avis_Rent_a_Car_System), accessed on Mar 20, 2014.




Figure 7: A 2D MDS similarity map of all the 18x2 rivals. Similarity measures are derived from their followers’ aggregated friends. Rival pairs are represented by stars of the same color. Some labels have been shortened due to space constraints. dunkinD represents DunkinDonuts and dem., Democrats.

## 6. WHEN BEING MAINSTREAM IS (NOT) BEST

Does being “mainstream” always equal being more popular? We try to answer this question by looking at whether Twitter accounts with the more *normal* followers typically also have *more* followers than their direct rivals. To establish what “normal” means we combine all the ~63,000 followers of the 18 rivals from the US that we used in our previous analysis. All following information for any of the seed accounts is removed and the remaining non-empty following feature vectors are normalized in 2-norm, and then summed and re-normalized. This “average centroid” is used as a reference for the follower preferences of a normal user. For each of the 18 pairs we then looked whether the rival whose followers were closest to this average also had more followers from the US. Here we estimated the number of followers from the US by first computing the fraction of users containing a profile location string as “US” out of a randomly sampled set of 10k users (see Section 3). This fraction was then applied to the total number of followers to obtain an estimate. Note that though the total number of followers from the US is likely larger, we followed this procedure to have an estimate about the user population from which the training/test users were drawn as for them the same procedure was applied. Only for three out of 18 cases the account whose followers were further away from the global average had more US followers than their rival. This gives statistically significant ( $p=.05$ ) evidence that, generally, mainstream means more followers. Related to the idea whether being mainstream pays off, we also look at whether having a homogeneous crowd of followers pays off. We quantify the degree of homogeneity by the average cosine similarity of two random followers of a particular seed account. This within-cluster similarity is shown in the last column of Table 5. For 14 out of 18 case the rival with the more homogeneous followership also had a bigger followership. We also note that the similarity scores between cluster centroids and the global average centroid are higher than the within-cluster similarity scores. This again indicates that a “global” approach works better for this data, as we saw in Section 4.1.

## 7. CONCLUSIONS

We presented an in-depth study of co-following behavior on Twitter which contributes to (i) a better understanding of language-agnostic user classification on Twitter, (ii) eliciting opportunities for Computational Social Science, and (iii) improving online marketing by identifying cross-selling opportunities. Concretely, we used the

Rival 1 vs. Rival 2	# fwers	Similarity	
		to aver. centroid	within cluster
@Budweiser vs. @MillerCoors	9,614 / 2,524	0.81 / 0.76	0.73 / 0.64
@FedEx vs. @UPS	27,681 / 21,585	0.78 / 0.78	0.76 / 0.71
@GM vs. @Ford	39,130 / 71,784	0.83 / 0.80	0.70 / 0.65
@GOP vs. @TheDemocrats	99,420 / 100,008	0.80 / 0.82	0.77 / 0.81
@Hertz vs. @Avis	10,523 / 2,214	0.84 / 0.83	0.68 / 0.51
@InsideFerrari vs. @lamborghini	19,304 / 19,422	0.81 / 0.85	0.63 / 0.69
@jcpenny vs. @Sears	78227 / 23133	0.85 / 0.83	0.78 / 0.77
@McDonalds vs. @BurgerKing	363,184 / 61,617	0.87 / 0.84	0.80 / 0.85
@MercedesBenz vs. @bmw	14,630 / 2,874	0.78 / 0.80	0.73 / 0.58
@Nike vs. @Reebok	276,747 / 34,259	0.81 / 0.80	0.76 / 0.79
@NikonUSA vs. @CanonUSAimaging	26,474 / 5,286	0.76 / 0.74	0.66 / 0.61
@pepsi vs. @CocaCola	240,481 / 127,723	0.83 / 0.78	0.81 / 0.75
@PUMA vs. @adidas	18,860 / 24,861	0.79 / 0.81	0.69 / 0.73
@SamsungMobile vs. @TheAppleInc	95,095 / 7,098	0.81 / 0.80	0.65 / 0.59
@Starbucks vs. @DunkinDonuts	922,533 / 142,709	0.87 / 0.83	0.66 / 0.79
@Target vs. @Walmart	265,424 / 125,778	0.81 / 0.78	0.78 / 0.75
@thewanted vs. @onedirection	111,166 / 693,313	0.81 / 0.80	0.88 / 0.75
@Visa vs. @MasterCard	42,723 / 18,780	0.86 / 0.84	0.81 / 0.76

Table 5: Results for the similarity analysis of the rivals’ centroids to the average centroid. The number of followers here is an estimate of the number of followers from US.

similarity of followers’ friends to predict a users’ preferences and to group main Twitter users according to their audiences’ similarities. We showed that such language-agnostic co-following information provides strong signals for diverse classification tasks and that these signals persist even when (i) the most discriminative features are removed and (ii) only relatively “sparse” users with fewer than 152 but more than 43 Twitter friends are considered. Rather than solely focusing on the classification task, we presented applications of our methodology to the area of Computational Social Science and confirmed stereotypes such as that the LGBT activist @ladygaga is more popular among @TheDemocrats followers than among @GOP followers. In the domain of marketing we gave evidence that celebrity endorsement is reflected in co-following and we demonstrated how our methodology can be used to reveal the audience similarities between Apple and Puma and, less obviously, between Nike and Coca-Cola. Concerning a user’s popularity we found a statistically significant connection between having a more “average” followership and having more followers than direct rivals. To the best of our knowledge, this is the first systematic study that shows how co-following on Twitter can be used for a variety of applications.

Our main focus in this paper was to introduce the concept of (2nd order) co-following and examine how it works for a wide range of settings, rather than the algorithm itself. In future, we would like to focus more on the algorithmic perspective and extend our work by looking deeper into aspects such as comparing how co-following fares in comparison to, e.g., methods which use the tweet content and user profile or compare the plots generated by MDS with other community detection algorithms.

## 8. REFERENCES

- [1] Republicans, democrats split on olive garden, other food issues, February 2013. <http://newsfeed.time.com/2013/02/27/republicans-democrats-split-on-olive-garden-other-food-issues/>.
- [2] J. An, M. Cha, P. K. Gummadi, and J. Crowcroft. Media landscape in twitter: A world of new conventions and political diversity. In *ICWSM*, 2011.
- [3] P. Barbera. Birds of the same feather tweet together: Bayesian ideal point estimation using twitter data. In *PolNet*, 2013. [https://files.nyu.edu/pba220/public/barbera\\_birds\\_july2013.pdf](https://files.nyu.edu/pba220/public/barbera_birds_july2013.pdf).
- [4] H. Bast and D. Majumdar. Why spectral retrieval works. In *SIGIR*, pages 11–18, 2005.
- [5] J. Bernstein, R. E. Bromley, and K. T. Meyer. Republicans and golf, democrats and outkast: Or, party political culture from the top down. *The Forum*, 4(3), 2006.
- [6] J. Blakley. The zogby/lear center survey on politics and entertainment, 2008. <http://www.learcenter.org/html/projects/?cm=zogby/08>.
- [7] C. Cioffi-Revilla. Computational social science. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2(3):259–271, 2010.
- [8] R. Cohen and D. Ruths. Classifying political orientation on twitter: It’s not easy! In *ICWSM*, 2013.
- [9] M. Conover, B. Gonçalves, J. Ratkiewicz, A. Flammini, and F. Menczer. Predicting the political alignment of twitter users. In *SocialCom/PASSAT*, pages 192–199, 2011.
- [10] M. J. Culnan. Mapping the intellectual structure of mis, 1980-1985: a co-citation analysis. *Mis Quarterly*, pages 341–353, 1987.
- [11] J. Davis and M. Goadrich. The relationship between precision-recall and roc curves. In *ICML*, pages 233–240, 2006.
- [12] P. Doyle. Favorite beers of republicans and democrats, October 2012. <http://www.bostonmagazine.com/news/blog/2012/10/01/favorite-beers-republicans-democrats/>.
- [13] S. Fortunato. Community detection in graphs. *Physics Reports*, 486(3-5):75–174, 2010.
- [14] L. Freeman. A set of measures of centrality based on betweenness. *Sociometry*, 40:35–41, 1977.
- [15] R. O. G. Gavilanes, D. Quercia, and A. Jaimes. Cultural dimensions in twitter: Time, individualism and power. In

- ICWSM, 2013.
- [16] S. Goel, J. M. Hofman, and M. I. Sirer. Who does what on the web: A large-scale study of browsing behavior. In *ICWSM*, 2012.
  - [17] S. A. Golder and S. Yardi. Structural predictors of tie formation in twitter: Transitivity and mutuality. In *SocialCom/PASSAT*, pages 88–95, 2010.
  - [18] B. A. Huberman, D. M. Romero, and F. Wu. Social networks that matter: Twitter under the microscope. *First Monday*, 14(1), 2009.
  - [19] C. Hutto, S. Yardi, and E. Gilbert. A longitudinal study of follow predictors on twitter. In *CHI*, pages 821–830, 2013.
  - [20] A. Java, X. Song, T. Finin, and B. L. Tseng. Why we twitter: understanding microblogging usage and communities. In *WebKDD/SNA-KDD*, pages 56–65, 2007.
  - [21] T. Kahute. Puma brand communication analysis, 2006. <http://www.scribd.com/doc/35130565/Puma-Brand-Analysis>.
  - [22] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *WWW*, pages 591–600, 2010.
  - [23] A. Lancichinetti and S. Fortunato. Community detection algorithms: A comparative analysis. *Phys. Rev. E*, 80:056117, Nov 2009.
  - [24] D. Lazer, A. S. Pentland, L. Adamic, S. Aral, A. L. Barabasi, D. Brewer, N. Christakis, N. Contractor, J. Fowler, M. Gutmann, et al. Life in the network: the coming age of computational social science. *Science*, 323(5915):721, 2009.
  - [25] M. E. J. Newman. Modularity and community structure in networks. *PNAS*, 103(23):8577–8582, 2006.
  - [26] Y. Pan, D.-H. Li, J.-G. Liu, and J.-Z. Liang. Detecting community structure in complex networks via node similarity. *Physica A: Statistical Mechanics and its Applications*, 389(14):2849–2857, 2010.
  - [27] M. Pennacchiotti and A.-M. Popescu. Democrats, republicans and starbucks aficionados: user classification in twitter. In *KDD*, pages 430–438, 2011.
  - [28] D. D. Posta, Y. Shi, and M. Macy. Why do liberals drink latte?, 2013. under review.
  - [29] D. Quercia, L. Capra, and J. Crowcroft. The social world of twitter: Topics, geography, and emotions. In *ICWSM*, 2012.
  - [30] D. Rao, D. Yarowsky, A. Shreevats, and M. Gupta. Classifying latent user attributes in twitter. In *SMUC*, pages 37–44, 2010.
  - [31] M. Shannon and W. Feltus. What your beer says about your politics, September 2012. <http://www.nationaljournal.com/blogs/hotlineoncall/2012/09/what-your-beer-says-about-your-politics-27>.
  - [32] H. Small. Co-citation in the scientific literature: A new measure of the relationship between two documents. *Journal of the American Society for information Science*, 24(4):265–269, 1973.
  - [33] H. Small. Co-citation in the scientific literature: A new measure of the relationship between two documents. *Journal of the American Society for information Science*, 24(4):265–269, 1973.
  - [34] L. Vaughan. Visualizing linguistic and cultural differences using web co-link data. *Journal of the American Society for Information Science and Technology*, 57(9):1178–1193, 2006.
  - [35] Y. Wang and M. Kitsuregawa. Link based clustering of web search results. In *Advances in Web-Age Information Management*, pages 225–236, 2001.
  - [36] I. Weber and A. Jaimes. Who uses web search for what: and how. In *WSDM*, pages 15–24, 2011.
  - [37] D. T. Wijaya and S. Bressan. Clustering web documents using co-citation, coupling, incoming, and outgoing hyperlinks: a comparative performance analysis of algorithms. *International Journal of Web Information Systems*, 2(2):69–76, 2006.
  - [38] W. Xie, C. Li, F. Zhu, E.-P. Lim, and X. Gong. When a friend in twitter is a friend in life. In *WebSci*, pages 344–347, 2012.
  - [39] D. Yin, L. Hong, and B. D. Davison. Structural link analysis and prediction in microblogs. In *CIKM*, pages 1163–1168, 2011.
  - [40] Z. Yin, M. Gupta, T. Weninger, and J. Han. A unified framework for link recommendation using random walks. In *ASONAM*, pages 152–159, 2010.
  - [41] F. A. Zamal, W. Liu, and D. Ruths. Homophily and latent attribute inference: Inferring latent attributes of twitter users from neighbors. In *ICWSM*, 2012.